arXiv:1406.6670v1 [math.ST] 25 Jun 2014

LEARNING THE ERGODIC DECOMPOSITION

NABIL I. AL-NAJJAR AND ERAN SHMAYA

Abstract. A Bayesian agent learns about the structure of a stationary process from observing past outcomes. We prove that his predictions about the near future become approximately those he would have made if he knew the long run empirical frequencies of the process.

1. Introduction

Consider a stationary, finite-valued stochastic process with probability law μ. According to the ergodic theorem, an observer of this process can reconstruct the 'true' ergodic component of the process from observing a single typical infinite realization. Decision problems, on the other hand, are often concerned with making predictions based on finite past observations. In such problems, the primary object of interest is the predictive distribution about the outcome of the process at a given day given the finite history of outcomes from previous days.

This paper relates these two perspectives on predictions and decisions. We consider the long-run properties of an observer's predictive distribution over next period's outcome as observations accumulate. We show that the predictive distribution becomes arbitrarily close to the predictive distribution conditioned on knowledge of the true ergodic component, in most periods almost surely. Thus, as data accumulates, an observer's predictive distributions based on finite history become the 'correct' predictions, in the sense of becoming nearly as good as what he would have predicted given knowledge of the objective empirical frequencies of the process. We demonstrate that the various qualifications we impose cannot be dropped.

Our results connect several literatures on learning and predictions in stochastic environments. First, there is the literature on the strong merging of opinions, pioneered by Blackwell and Dubins [2].[1] More directly relevant to our purpose are the weaker notions of merging introduced by Kalai and Lehrer [9] and Lehrer and Smorodinsky [11], which focus on closeness of near-horizon predictive distributions. While strong merging obtains only under stringent assumptions, weak merging can be more easily satisfied. In our setting, for example, the

Date: First draft: February 2013; This version: June 26, 2014.
2000 *Mathematics Subject Classification.* Primary: 60G10, 91A26. Secondary: 37A25,62M20, 62F15.
We thank Ehud Kalai, Ehud Lehrer and Rann Smorodinsky for helpful discussions.

[1] Kalai and Lehrer [7] apply this concept to learning in games.

posteriors may fail to strongly merge with the true parameter, no matter how much data accumulates. This strong notion of merging is unnecessary in contexts where decision makers discount the future or care only about a fixed number of future periods. Weak merging, to which our results apply, is usually sufficient.

Another line of enquiry focuses on representations of the form $\mu = \int \mu_\theta \, \mathrm{d}\lambda(\theta)$, where a probability measure μ (the law of the stochastic process) is expressed as a convex combination of distributions $\{\mu_\theta\}_{\theta\in\Theta}$ that may be viewed as especially "simple," or "elementary." Such representations, also called *decompositions*, are useful in models of learning where the set of parameters Θ may be viewed as the main object of learning. Two seminal theorems are de Finetti's representation of exchangeable distributions and the ergodic decomposition theorem for stationary processes. Exchangeability rules out many interesting patterns of inter-temporal correlation, so it is natural to consider the larger class of stationary distributions. For this class, the canonical decomposition is in terms of the ergodic distributions. This is the finest decomposition possible using parameters that are themselves stationary. Our main theorem states that a Bayesian decision maker's predictions, based on finite histories, become arbitrarily close to those he would have made given knowledge of the true ergodic component.

Our result should also be contrasted with Doob's consistency theorem which states that Bayesian posteriors weakly converge to the true parameter. When the focus is the quality of decisions, what matters is not the agents' belief about the true parameter but the quality of his predictions. Although the two concepts are related, they are not the same. The difference is seen in the following example from Jackson, Kalai and Smorodinsky [6, Example 5]: Assume that the outcomes Heads and Tails are generated by tossing a fair coin. If we take the set of all dirac measures on infinite sequences of Heads-Tails outcomes as "parameters", then the posterior about the parameter converges weakly to a belief that is concentrated on the true realization. On the other hand the agent's predictions about next period's outcome is constant and never approach the predictions given the true "parameter." This example highlights that convergence of posterior beliefs to the true parameters may have little relevance to an agent's predictions and behavior.

Every process can be represented in an infinite number of ways, many of which, like the decomposition of the coin toss process above, are not very sensible. Jackson, Kalai and Smorodinsky [6] study the question of what makes a particular decomposition of a stochastic process sensible. One requirement is for the process to be learnable, in the sense that an agent's predictions about near-horizon events become close to what he would have predicted had he known the true parameter. Given the close connection between ergodic distributions and long-run frequencies, the most natural decomposition $\mu = \int \mu_\theta \, \mathrm{d}\lambda(\theta)$ of a stationary process is where the θ's index the ergodic distributions. We show that their results do not

apply to the class of stationary processes and their canonical ergodic decompositions. We show, however, that the ergodic decomposition is learnable in a weaker, yet meaningful sense as described below.

A third related literature, which traces to Cover [3], is non-Bayesian estimation of stationary processes. See Morvai and Weiss [13] and the reference therein. This literature looks for an algorithm that make near-horizon predictions that are accurate for every stationary process. Our proofs of Theorem 3.1 and Example 3.3 rely on techniques that were developed in this literature. There is however a major difference between that literature and our work: We are interested in a specific algorithm, namely Bayesian updating. Our agent's predictions and behavior are derived from this updating process. We show how to apply the mathematical apparatus developed for the non-Bayesian estimation in our Bayesian setup.

2. Formal model

2.1. **Preliminaries.** An agent (a decision maker, a player, or a statistician) observes a stochastic process $(\zeta_0, \zeta_1, \zeta_2, \dots)$ that takes values in a finite set of *outcomes* A. Time is indexed by n and the agent starts observing the process at $n = 0$. Let $\Omega = A^{\mathbb{N}}$ be the space of realizations of the process, with generic element denoted $\omega = (a_0, a_1, \dots)$. Endow Ω with the product topology and the induced Borel structure $\mathcal{F}$. Let $\Delta(\Omega)$ be the set of probability distributions over Ω. The law of the process is an element μ of $\Delta(\Omega)$. A standard way to represent uncertainty about the process is in terms of an index set of "parameters:"

Definition 2.1. Let $\mu \in \Delta(\Omega)$. A *decomposition* of μ is a quadruple $(\Theta, \mathcal{B}, \lambda, (\mu_\theta))$ where: $(\Theta, \mathcal{B}, \lambda)$ is a standard probability space of *parameters* and $\mu_\theta \in \Delta(\Omega)$ for every $\theta \in \Theta$ such that the map $\theta \mapsto \mu_\theta(A)$ is $\mathcal{B}$-measurable and

$$\mu(A) = \int_\Theta \mu_\theta(A)\ \lambda(\mathrm{d}\theta) \tag{1}$$

for every $A \in \mathcal{F}$. ▲

A decomposition captures a certain way in which a Bayesian agent arranges his beliefs: The agent views the process as a two stages randomization. First a parameter θ is chosen according to λ and then the outcomes are generated according to μ_θ. Beliefs can be represented in many ways. The two extreme decompositions are: (1) the *Trivial Decomposition.* with $\Theta = \{\bar{\theta}\}$, $\mathcal{B}$ is trivial, and $\mu_{\bar{\theta}} = \mu$; and (2) the *Dirac Decomposition.* with $\Theta = A^{\mathbb{N}}$, $\mathcal{B} = \mathcal{F}$, and $\lambda = \mu$. A "parameter" in this case is just a Dirac measure δ_ω that assigns probability 1 to the realization ω.

We are interested in decompositions that identify "useful" patterns shared by many realizations. These patterns capture our intuition of fundamental properties of a process. The two extreme cases are usually unsatisfactory. In the Dirac decomposition, there are as many

parameters as there are realizations; parameters simply copy realizations. In the trivial decomposition, there is a single parameter and thus cannot discriminate between different interesting patterns.

Stationary beliefs admit a well-known decomposition with natural properties. Recall that the set of stationary measures over Ω is convex and compact in the weak*-topology. Its extreme points are called *ergodic beliefs*. We denote the set of ergodic beliefs by $\mathcal{E}$. Every stationary belief μ admits a unique decomposition in which the parameter set is the set of ergodic beliefs: $\mu = \int \nu \ \lambda(\mathrm{d}\nu)$ for some belief $\lambda \in \Delta(\mathcal{E})$. We call this decomposition *the ergodic decomposition*.

According to the ergodic theorem, for every stationary belief μ and every block $(\bar{a}_0, \dots, \bar{a}_{k-1}) \in A^k$, the limit frequency

$$\Pi(\omega; \bar{a}_0, \dots, \bar{a}_{k-1}) = \lim_{n\to\infty} \frac{1}{n} \#\big\{0 \le t < n : a_t = \bar{a}_0, \dots, a_{t+k-1} = \bar{a}_{k-1}\big\}$$

exists for μ-almost every realization $\omega = (a_0, a_1, \dots)$. When μ is ergodic this limit equals the probability $\mu([\bar{a}_0, \dots, \bar{a}_{k-1}])$. Thus, for ergodic processes, the probability of every block equals its (objective) empirical frequency.

The ergodic decomposition theorem states that for μ-almost every ω, The function $\Pi(\omega; \cdot)$ defined over blocks can be extended to a stationary measure over $\Delta(\Omega)$ which is also ergodic. Moreover, $\mu = \int \Pi(\omega; \cdot) \ \mu(\mathrm{d}\omega)$, so that the function $\omega \to \Pi(\omega; \cdot)$ recovers the ergodic parameter from the realization of the process. Thus, the parameters μ_θ in the ergodic decomposition represent the empirical distribution of finite sequences of outcome along the realization of the stationary process. These parameters capture our intuition of fundamentals of the process.

A special case of the ergodic decomposition is the decomposition of an exchangeable distribution μ via i.i.d. distributions. For future reference, consider the following example:

Example 2.2. The set of outcomes is $A = \{0, 1\}$ and the agent's belief is given by

$$\mu\left(\zeta_n = a_0, \dots, \zeta_{n+k-1} = a_{k-1}\right) = \frac{1}{(k+1) \cdot \binom{k}{d}}$$

for every $n, k \in \mathbb{N}$ and $a_0, \dots, a_{k-1} \in A$ where $d = a_0 + \cdots + a_{k-1}$. Thus, the agent believes that if he observes the process k consecutive periods then the number d of good periods (periods with outcome 1) is distributed uniformly in $[0, k]$ and all configuration with d good outcomes are equally likely.

De-Finetti's decomposition is given by $(\Theta, \mathcal{B}, \lambda)$ where $\Theta = [0,1]$ equipped with the standard Borel structure $\mathcal{B}$ and Lebesgue's measure λ, and, for $\theta \in \Theta$ $\mu_\theta \in \Delta(\Omega)$ is the distribution of i.i.d coin tosses with probability of success θ:

$$\mu_\theta\left(\zeta_n = a_0, \ldots, \zeta_{n+k-1} = a_k\right) = \theta^d(1-\theta)^{k-d}$$

▲

2.2. **Learning.** For every $\mu \in \Delta(\Omega)$ and sequence $(a_0, \ldots, a_{n-1}) \in A^n$ with positive μ-probability, the *n-period predictive distribution* is the element $\mu(\cdot|a_0, \ldots, a_{n-1}) \in \Delta(A)$ representing the agent's prediction about next period's outcomes given a prior μ and after observing the first n outcomes of the process. Predictive distributions in this paper will always refer to one-step ahead predictions. This is for expository simplicity; our analysis covers any finite horizon.

Kalai and Lehrer [9], and Kalai, Lehrer and Smorodinsky [8] introduced the following notions merging. Note that in our setup, where the set of outcomes is the same in every period, this definition of merging is the same as 'weak star merging' in D'Aristotile, Diaconis and Freedman [4].

Definition 2.3. Let $\mu, \tilde{\mu} \in \Delta(\Omega)$. Then the belief $\tilde{\mu}$ *merges* to μ if

$$|\tilde{\mu}(\cdot|a_0, \ldots, a_{n-1}) - \mu(\cdot|a_0, \ldots, a_{n-1})| \xrightarrow[n\to\infty]{} 0$$

for μ-almost every realization $\omega = (a_0, a_1, \ldots) \in A^{\mathbb{N}}$.

The belief $\tilde{\mu}$ *weakly merges* to μ if

$$|\tilde{\mu}(\cdot|a_0, \ldots, a_{n-1}) - \mu(\cdot|a_0, \ldots, a_{n-1})| \xrightarrow[n\to\infty]{\text{s.c}} 0.^{2} \tag{2}$$

for μ-almost every realization $\omega = (a_0, a_1, \ldots) \in A^{\mathbb{N}}$. ▲

Here and later, for every pair $p, q \in \Delta(A)$ we let $\|p - q\| = \max_{a\in A} |p[a] - q[a]|$. These definitions were inspired by Blackwell and Dubins idea of strong merging, which requires that the prediction of $\tilde{\mu}$ will be similar to the prediction of μ not just for the next period but for the infinite horizon.

Definition 2.4. A decomposition $(\Theta, \mathcal{B}, \lambda, (\mu_\theta))$ of $\mu \in \Delta(\Omega)$ is *learnable* if μ merges with μ_θ for λ-almost every θ. The decomposition is *weakly learnable* if μ weakly merges with μ_θ for λ-almost every θ. ▲

[2] A bounded sequence of real numbers $a_0, a_1, \ldots$ is said to *strongly Cesaro converges* to a real number a, denoted $a_n \xrightarrow[n\to\infty]{\text{s.c}} a$, if $\lim_{n\to\infty} \frac{1}{n} \sum_{k=0}^{n-1} |a_k - a| = 0$. Equivalently, $a_n \xrightarrow{\text{s.c}} a$ if there exists a full density set $T \subseteq \mathbb{N}$ such that $\lim_{n\to\infty, n\in T} a_n = a$.

As an example of a learnable decomposition, consider the Bayesian agent of Example 2.2. In this case

$$\mu(1|a_0,\dots,a_{n-1}) = \frac{a_0+\cdots+a_{n-1}+1}{n+1}.$$

For every $\theta \in [0,1]$ the strong law of large numbers implies that for every parameter $\theta \in [0,1]$ this expression converges μ_θ-almost surely to θ. Therefore μ merges with μ_θ for every θ, so De Finetti's decomposition is learnable (and, a fortiori, weakly learnable). This is a rare case in which the predictions $\mu(\zeta_n \in \cdot|a_0,\dots,a_{n-1})$ and $\mu_\theta(\zeta_n \in \cdot|a_0,\dots,a_{n-1})$ can be calculated explicitly. In general merging and weak merging are difficult to establish, because the Bayesian prediction about the next period is a complicated expression which is potentially depends on entire observed past.

2.3. **Motivation for Weak Merging.** In applications, μ represent the true process generating observations, and $\tilde{\mu}$ is a Bayesian agent's belief. To say that $\tilde{\mu}$ weakly merges with μ means that his next period predictions are accurate except for rare times.

To connect this concept with statistical decision problems, suppose that in every period, before the outcome is realized, the agent has to take some *decision* from a finite set D. The agent's *payoff* is represented by the *payoff function* $r : A \times D \to [0,1]$. A *strategy* is given by $f : \bigcup_{n\geq 0} A^n \mapsto D$, with $f(a_0,\dots,a_{n-1})$ denoting the action taken given the past realized outcomes. Let

$$V_N(f) = \frac{1}{N}\int \sum_{n=0}^{N-1} r\left(a_n, f(a_0,\dots,a_{n-1})\right)\ \mathrm{d}\mu$$

be the expected average expected payoff in the first N periods. Fix $\epsilon > 0$. A strategy f^* is *ϵ-optimal for N periods under μ* if $V_N(f) \leq V_N(f^*) + \epsilon$ for every strategy f. Of course the optimal strategy depends on the agent's belief μ. The following proposition, which is immediate from the definition of weak learning, says that an agent who maximizes according to a belief that weakly merges with the truth will play ϵ-optimal strategies against the truth if he is sufficiently patient. By 'sufficiently patient' we mean that the horizon N is large. Similar result applies if the agent aggregates periods' payoffs using some discount factor where by 'sufficiently patient' is meant that the discount factor is close to 1.

Proposition 2.5. *Let $\tilde{\mu}, \mu \in \Delta(\Omega)$ be such that $\tilde{\mu}$ weakly merges with μ. For every $\epsilon > 0$ there exists N_0 such that for every $N > N_0$, in every decision problem, every 0-optimal strategy for N periods under $\tilde{\mu}$ is ϵ-optimal for N periods under μ.*

Kalai, Leher and Smorodinsky [8] provide a motivation for weak learning in terms of the properties of calibration tests. The idea of calibration originated with Dawid [5]. A calibration test of a forecast compares the predicted frequency of events to their realized empirical frequencies. Kalai et al. showed that $\tilde{\mu}$ weakly merges with μ if and only if

forecasts made by $\tilde{\mu}$ pass all calibration tests of a certain type when the outcomes are generated according to μ.

Finally, Lehrer and Smorodinsky [12] provide a characterization of weak merging in terms of the relative entropy between $\tilde{\mu}$ and μ.[3] No similar characterization is known for merging.

2.4. **Merging and the Consistency of Bayesian Estimators.** The idea of learning captured by Definition 2.4 concerns the quality of predictions made about near-horizon events. Another, perhaps more common, way to think about Bayesian inference is in terms of the consistency of Bayesian estimator. Consistency can be thought of as concerning learning the parameter itself. Recall that the Bayesian estimator of the parameter θ is the agent's conditional belief over θ after observing the outcomes of the process. It is well known that under any 'reasonable' decomposition, the Bayesian estimator is consistent, i.e., the estimator weakly converges to the Dirac measure over the true parameter as data accumulates[4]. However, consistency of the estimator does not imply that the agent can use what he has learned to make predictions about future outcomes. For example, consider the Dirac decomposition of the process of fair coin tosses. Suppose the true parameter is ω^* for some $\omega^* = (\omega_0^*, \omega_1^*, \dots)$. After observing the first n outcomes of the process the agent's belief about the parameter is uniform over all ω that agrees with ω^* on the first n coordinates. While this belief indeed converges to δ_{ω^*}, the agent does not gain any new insight about the future of the process from learning the parameter. This decomposition is therefore not learnable.

3. Main Theorem

We are now in a position to state our main theorem.

Theorem 3.1. *The ergodic decomposition of every stationary stochastic process is weakly learnable.*

To see the implications of our theorem, consider the following Hidden Markov process

Example 3.2. An agent believes that the state of the economy every period is a noisy signal of an underlying "hidden" states that changes according to a Markov chain with memory 1. Formally, let $A = \{\mathrm{B}, \mathrm{G}\}$ be the set of outcomes, $H = \{\mathrm{B}, \mathrm{G}\}$ the set of hidden (unobserved) states, and (ξ_n, ζ_n) a $(H \times A)$-valued stationary Markov process with transition matrix $\rho : H \times A \to \Delta(H \times A)$ given by

$$\rho(h, a)[h', a'] = \big(p\delta_{h,h'} + (1-p)(1-\delta_{h,h'})\big) \cdot \big(q\delta_{h',a'} + (1-q)(1-\delta_{h',a'})\big),$$

[3] However, we do not know whether their condition can be used to prove our theorem without repeating the whole argument.

[4] The argument traces back to Doob. See, for example, Weizsacker [16] and the references therein. It holds whenever the decomposition has the property that the realization of the process determines the parameter

where $1/2 < p, q < 1$. Thus, if the hidden state in period n was h then at period $n+1$ the hidden state h' remains h with probability p and changes with probability $1-p$. The observed state a' of period $n+1$ will then be h' with probability q and is different from h with probability $1-q$. Let $\mu_{p,q} \in \Delta(A^{\mathbb{N}})$ be the distribution of $\zeta_0, \zeta_1, \ldots$. Then $\mu_{p,q}$ is a stationary process which is not markov of any order. If the agent is uncertain about p, q then his belief μ about the outcome process is again stationary, and can be represented by some prior over the parameter set $\Theta = (1/2, 1] \times (1/2, 1]$. This decomposition of μ will be the ergodic decomposition. ▲

The consistency of the Bayesian estimator for (p, q) implies that the conditional belief over the parameter (p, q) converges almost surely in the weak-topology over $\Delta(\Theta)$ to the belief concentrated on the true parameter. However, because next-period's predictions involve complicated expressions that depend on the entire history of the process, it is not clear whether these predictions merge with the truth. It follows from our theorem that they weakly merge.

Consider now the general case. If the agent knew the fundamental θ, then at period n, after observing the partial history $(a_0, \ldots, a_{n-1})$, his predictive probability that the next period outcome is a_n would have been

$$\frac{\mu_\theta(a_0, \ldots, a_{n-1}, a_n)}{\mu_\theta(a_0, \ldots, a_{n-1})}. \tag{3}$$

Again consistency of the Bayesian estimator implies that, given uncertainty about the fundamental, the agent's assessment of $\mu_\theta(b)$ becomes asymptotically accurate for every block b. However, when the agent has to compute the next-period posterior probability (3), he only had one observation of a block of size n and no observation of the block of size $n+1$ so at that stage his assessment of the probabilities that appear in (3) may be completely wrong. Our theorem says that the agent would still weakly learn to make these predictions correctly.

Theorem 3.1 states that the agent will make predictions about near-horizon events as if he knew the fundamental of the process. Note, however, that it is not possible to ensure that the agent will learn to predict long-run events correctly, no matter how much data accumulates. For example, consider an agent who faces a sequence of i.i.d. coin tosses with parameter $\theta \in [0, 1]$ representing the probability of Heads. Suppose this agent has a uniform prior over [0,1]. This agent will eventually learn to predict near horizon outcomes as if he knew the true parameter θ, but if he will continue to assign probability 0 to the event that the long-run frequency is θ. In economic models, discounting implies that only near-horizon events matter.

We end this section with an example that in Theorem 3.1 weak learnability cannot be replaced by learnability. The example is a modification of an example given by Ryabko for the forward prediction problem in a non-Bayesian setup [14].

Example 3.3. Every period there is a probability 1/2 for eruption of war. If no war erupts then the outcome is either bad economy or good economy and is a function of the number of peaceful periods since the last war. The function from the number of peaceful periods to outcome is an unknown parameter of the process, and the agent has a uniform prior over this parameter.

Formally, let $A = \{\mathrm{W}, \mathrm{B}, \mathrm{G}\}$ be the set of outcomes. We define $\mu \in \Delta(A^{\mathbb{N}})$ through its ergodic decompositions. Let $\Theta = \{\mathrm{B}, \mathrm{G}\}^{\{1,2,\dots\}}$ be the set of parameters with the standard Borel structure $\mathcal{B}$ and the uniform distribution λ. Thus, a parameter is a function $\theta : \{1, 2, \dots\} \to \{\mathrm{B}, \mathrm{G}\}$. We can think about this belief as a hidden markov model where the unobservable process $\xi_0, \xi_1, \dots$ is the time that elapsed since last time a war occurred. Thus, $\xi_0, \xi_1, \dots$ is the $\mathbb{N}$-valued stationary Markov process with transition probability

$$\rho(j|k) = \begin{cases} 1/2, \text{ if } j = k+1, \\ 1/2, \text{ if } j = 0, \\ 0, \text{ otherwise.} \end{cases}$$

for every $j, k \in \mathbb{N}$, and μ_θ is the distribution of a sequence $\zeta_0, \zeta_1, \dots$ of A-valued random variables such that

$$\zeta_n = \begin{cases} \mathrm{W}, & \text{if } \xi_n = 0 \\ \theta(\xi_n), & \text{otherwise.} \end{cases}$$

Consider a Bayesian agent who observes the process. After the first time a war erupts the agent keeps track of the state of the process ξ_n at every period. If there is no uncertainty about the parameter, i.e., if the Bayesian agent knew θ, his prediction about the next outcome when $\zeta_n = k$ gives probability 1/2 to outcome W and probability 1/2 to outcome $\theta(k+1)$. On the other hand, if the agent does not know θ but believes that it is randomized according to λ, he can deduce the values $\theta(k)$ gradually while he observes the process. However for every $k \in \{1, 2, 3, \dots\}$ there will be a time when the agent will observe k consecutive peaceful period for the first time and at this point the agent's prediction about the next outcome will be $(1/2, 1/4, 1/4)$. Thus there will always be infinitely many occasions in which an agent who predicts according to μ will differ than an agent who predicts according to μ_θ. Therefore the decomposition is not learnable. On the other hand, in agreement with our theorem, these occasions become more infrequent as time goes by so the decomposition is weakly learnable. ▲

4. Proof of Theorem 3.1

Up to now we assumed that the stochastic process starts at time $n = 0$. When working stationary processes it is natural to extend the index set of the process from $\mathbb{N}$ to $\mathbb{Z}$, i.e. to assume that the process has infinite past. This is without loss of generality: every stationary stochastic process $\zeta_0, \zeta_1, \dots$ admits an extension $\dots, \zeta_{-1}, \zeta_0, \zeta_1, \dots$ to the index set $\mathbb{Z}$ [10, Lemma 10.2]. We therefore assume hereafter, with harmless contrast with our previous notation, that $\Omega = A^{\mathbb{Z}}$.

Let $\mathcal{D}$ be a σ-algebra Borel subsets of Ω. The *quotient space* of $(\Omega, \mathcal{F}, \mu)$ with respect to $\mathcal{D}$ is the unique (up to isomorphism of measure spaces) standard probability space $(\Theta, \mathcal{B}, \lambda)$ and a measurable map $\alpha : \Omega \to \Theta$ such that $\mathcal{D}$ is generated by α, i.e., for every $\mathcal{F}$-measurable function f from Ω to some standard probability space there exists a (unique up to equality λ-almost surely) $\mathcal{B}$-measurable *lifting* $\tilde{f}$ defined over Θ such that $f = \tilde{f} \circ \alpha \quad \mu - \text{a.s.}$. The *conditional distributions* of μ over $\mathcal{D}$ is the unique (up to equality λ-almost surely) family μ_θ of probability measures over $(\Omega, \mathcal{F}, \mu)$ such that:

(1) For every $\theta \in \Theta$ it holds that

$$\mu_\theta\left(\{\omega | \alpha(\omega) = \theta\}\right) = 1. \tag{4}$$

(2) The map $\theta \mapsto \mu_\theta(A)$ is $\mathcal{B}$-measurable and (1) is satisfied for every $A \in \mathcal{F}$.

We call $(\Theta, \mathcal{B}, \lambda, \mu_\theta)$ the *decomposition of μ induced by $\mathcal{D}$*. For every belief $\mu \in \Delta(\Omega)$, the trivial decomposition of μ is generated by the trivial sigma-algebra $\{\emptyset, \Omega\}$, the Dirac decomposition is generated by the sigma-algebra of all Borel subsets of Ω. The ergodic decomposition is induced by the σ-algebra $\mathcal{I}$ of all *invariant* Borel sets of Ω, i.e. all Borel sets $S \subseteq \Omega$ such that $S = T^{-1}(S)$ where $T : \Omega \to \Omega$ is the left shift.

We will prove a more general theorem, which may be interesting in its own right. Let $T : A^{\mathbb{Z}} \to A^{\mathbb{Z}}$ be the left shift so that $T(\omega)_n = \omega_{n+1}$ for every $n \in \mathbb{Z}$. A sigma-algebra $\mathcal{D}$ of Borel subsets of Ω is *shift-invariant* if $S \in \mathcal{D} \leftrightarrow T(S) \in \mathcal{D}$ for every Borel subset S of $A^{\mathbb{Z}}$.

Theorem 4.1. *Let μ be a stationary distribution over Ω and let $\mathcal{D}$ be a shift invariant σ-algebra of subsets of Ω such that $\mathcal{D} \subseteq \mathcal{F}^0_{-\infty}$. Then the decomposition of μ induced by $\mathcal{D}$ is weakly learnable.*

Theorem 3.1 follows immediately from Theorem 4.1 since the sigma-algebra of invariant sets $\mathcal{I}$ which induces the ergodic decomposition satisfies the assumption of the Theorem 4.1. We will prove Theorem 4.1 using Lemma 4.2

Lemma 4.2. *Let μ be a stationary distribution over $A^{\mathbb{Z}}$ and let $\mathcal{D}$ be a shift invariant σ-algebra of Borel subsets of $A^{\mathbb{Z}}$. Then*

$$\|\mu(\zeta_n = \cdot | \mathcal{F}^n_0 \vee \mathcal{D}) - \mu(\zeta_n = \cdot | \mathcal{F}^n_{-\infty} \vee \mathcal{D})\| \xrightarrow[n\to\infty]{s.c} 0 \ \mu\text{-a.s.} \tag{5}$$

Consider the case in which $\mathcal{D} = \{\emptyset, \Omega\}$ is trivial. Then Lemma 4.2 says that a Bayesian agent who observes a stationary process from time $n = 0$ onwards will make predictions in the long run as if he knew the infinite history of the process.

Proof of Lemma 4.2. For every $n \geq 0$ let $f_n : \Omega \to \Delta(A)$ be a version of the conditional distribution of ζ_0 according to μ given the finite history $\zeta_{-1}, \ldots, \zeta_{-n}$ and $\mathcal{D}$:

$$f_n = \mu(\zeta_0 = \cdot | \mathcal{F}^0_{-n} \vee \mathcal{D}),$$

and let $f_\infty : \Omega \to \Delta(A)$ be a version of the conditional distribution of ζ_0 according to μ given the infinite history $\zeta_{-1}, \ldots$ and $\mathcal{D}$:

$$f_\infty = \mu(\zeta_0 = \cdot | \mathcal{F}^0_{-\infty} \vee \mathcal{D}).$$

Let $g_n = \|f_n - f_\infty\|$. By the martingale convergence theorem $\lim_{n\to\infty} f_n = f_\infty$ μ-a.s and therefore

$$\lim_{n\to\infty} g_n = 0 \ \mu\text{-a.s} \tag{6}$$

It follows from the stationarity of μ and the fact that $\mathcal{D}$ is shift invariant that

$$\|\mu\left(\zeta_n = \cdot | \mathcal{F}^n_0 \vee \mathcal{D}\right) - \mu\left(\zeta_n = \cdot | \mathcal{F}^n_{-\infty} \vee \mathcal{D}\right)\| = \|f_n \circ T^n - f_\infty \circ T^n\| = g_n \circ T^n \ \mu\text{-a.s} \tag{7}$$

Therefore

$$\frac{1}{N}\sum_{n=0}^{N-1} \|\mu\left(\zeta_n = \cdot | \mathcal{F}^n_0 \vee \mathcal{D}\right) - \mu\left(\zeta_n = \cdot | \mathcal{F}^n_{-\infty} \vee \mathcal{D}\right)\| = \frac{1}{N}\sum_{n=0}^{N-1} g_n \circ T^n \xrightarrow[N\to\infty]{} 0 \ \mu\text{-a.s}$$

where the equality follows from (7) and the limit follows from (6) and Maker's generalization of the ergodic theorem [10, Corollary 10.8] to cover multiple functions simultaneously:

Maker's Ergodic Theorem. *Let $\mu \in \Delta(\Omega)$ be such that $T\mu = \mu$ and let $g_0, g_1, \cdots : \Omega \to \mathbf{R}$ be such that $\sup_n |g_n| \in L^1(\mu)$ and $g_n \to g_\infty \quad \mu - a.s.$ Then*

$$\frac{1}{N}\sum_{n=0}^{N-1} g_n \cdot T^n \xrightarrow[N\to\infty]{} E(g_\infty | \mathcal{I}) \quad \mu - a.s.$$

□

Proof of Theorem 4.1. From $\mathcal{D} \subseteq \mathcal{F}^0_{-\infty}$ it follows that $\mathcal{F}^n_{-\infty} \vee \mathcal{D} = \mathcal{F}^n_{-\infty}$. Therefore, from Lemma 4.2 we get that

$$\|\mu(\zeta_n = \cdot | \mathcal{F}^n_0 \vee \mathcal{D}) - \mu(\zeta_n = \cdot | \mathcal{F}^n_{-\infty})\| \xrightarrow[n\to\infty]{\text{s.c}} 0 \ \mu\text{-a.s.}$$

By the same lemma (with $\mathcal{D}$ trivial)

$$\|\mu(\zeta_n = \cdot | \mathcal{F}^n_0) - \mu(\zeta_n = \cdot | \mathcal{F}^n_{-\infty})\| \xrightarrow[n\to\infty]{\text{s.c}} 0 \ \mu\text{-a.s.}$$

By the last two limits and the triangular inequality

$$\|\mu(\zeta_n = \cdot|\mathcal{F}_0^n) - \mu(\zeta_n = \cdot|\mathcal{F}_0^n \vee \mathcal{D})\| \leq \tag{8}$$
$$\|\mu(\zeta_n = \cdot|\mathcal{F}_0^n) - \mu(\zeta_n = \cdot|\mathcal{F}_{-\infty}^n)\| + \|\mu(\zeta_n = \cdot|\mathcal{F}_0^n \vee \mathcal{D}) - \mu(\zeta_n = \cdot|\mathcal{F}_{-\infty}^n)\| \xrightarrow[n\to\infty]{\text{s.c}} 0 \ \mu\text{-a.s}$$

Let $(\Theta, \mathcal{B}, \lambda)$ be the quotient of $(\Omega, \mathcal{F}, \mu)$ over $\mathcal{D}$ and let (μ_θ) be the corresponding conditional distributions. Let S be the set of all realizations $\omega = (\dots, a_{-1}, a_0, a_1, \dots)$ such that

$$\|\mu(\zeta_n = \cdot|a_{n-1}, \dots, a_0) - \mu_\omega(\zeta_n = \cdot|a_{n-1}, \dots, a_0)\| \xrightarrow[n\to\infty]{\text{s.c}} 0.$$

Then $\mu(S) = 1$ by (8). But $\mu(S) = \int \mu_\theta(S)\lambda(\mathrm{d}\theta)$. It follows that $\mu_\theta(S) = 1$ for λ-almost every θ, a desired. □

5. Ergodicity and mixing

Mixing conditions formalize the intuition that observing a sequence of outcomes of a process does not change one's belief about events in the far future. Standard examples of mixing processes are i.i.d. processes and non-periodic markov processes. In this section we recall a mixing condition that was called "sufficiency for prediction" in JKS, show that the ergodic decomposition is not necessarily sufficient for prediction and show that a finer decomposition than the ergodic decomposition is sufficient for prediction and also weakly learnable.

Let $\overrightarrow{\mathcal{T}} = \bigwedge_{m\geq 0} \mathcal{F}_m^\infty$ be the *future tail* sigma-algebra where $\mathcal{F}_m^\infty$ the σ-algebra of Ω that is generated by $(\zeta_m, \zeta_{m+1}, \dots)$. A probability distribution (not necessarily stationary) $\nu \in \Delta(\Omega)$ is *mixing* if it is $\overrightarrow{\mathcal{T}}$-trivial, i.e., if $\nu(B) \in \{0, 1\}$ for every $B \in \overrightarrow{\mathcal{T}}$.[5]

If we want the components of the decomposition to be mixing we need a finer decomposition than the ergodic decomposition. This decomposition is the decomposition that is induced by the tail $\overrightarrow{\mathcal{T}}$ as shown in the following proposition.

Proposition 5.1. *Let $(\Theta, \mathcal{B}, \lambda, (\mu_\theta))$ be the decomposition of a belief $\mu \in \Delta(\Omega)$ that is induced by the tail $\overrightarrow{\mathcal{T}}$. Then μ_θ is mixing for λ-almost every θ.*

Proof. This proposition is JKS' Theorem 1. We repeat the argument here to clarify a gap in their proof.

The proposition follows from the fact that the conditional distributions of every probability distribution $\mu \in \Delta(\Omega)$ over the tail are almost surely tail-trivial (i.e., mixing). This fact was

[5] An equivalent way to write this condition is that for every n and ϵ, there is m such that

$$|\nu(B|a_0, \dots, a_{n-1}) - \nu(B)| < \epsilon$$

for every $B \in \mathcal{F}_m^\infty$ and partial history $(a_0, \dots, a_{n-1}) \in A^n$. JKS call such belief *sufficient for prediction*. They establish the equivalence with the mixing condition in their proof of their Theorem 1

recently proved by Berti and Rigo [1, Theorem 15][6]. We note that it is not true for every sigma-algebra $\mathcal{D}$ that the conditional distributions of μ over $\mathcal{D}$ are almost surely $\mathcal{D}$-trivial. This property is very intuitive (and indeed, easy to prove) when $\mathcal{D}$ is generated by a finite partition, or more generally when $\mathcal{D}$ is countably generated, but the tail is not countably generated, which is why Berti and Rigo's result is required. □

The next theorem uses Lemma 4.2 to show that the tail decomposition is also weakly learnable. In particular, Theorem 5.2 implies that the ergodic decomposition does not capture all the learnable properties of a stationary process.

Theorem 5.2. *The tail decomposition of a stationary stochastic process is weakly learnable.*

Proof. From Lemma 4.2 it follows that the decomposition induced by the past tail $\overleftarrow{\mathcal{T}}$ is learnable, since the past tail is shift invariant.

The theorem now follows from the fact that for every stationary belief μ over a finite set of outcomes it holds that $\overleftarrow{\mathcal{T}}_\mu = \overrightarrow{\mathcal{T}}_\mu$ where $\overleftarrow{\mathcal{T}}_\mu$ and $\overrightarrow{\mathcal{T}}_\mu$ are the completions of the past and future tails under μ. See Weiss [15, Section 7]. Therefore, the decomposition of μ induced by $\overrightarrow{\mathcal{T}}$ equals the decomposition induced by $\overleftarrow{\mathcal{T}}$, which is learnable. We note that the equality of the past and future tails of a stationary process is not trivial, it relies on finiteness of the set of outcomes A, and the proof relies on the notion of entropy. □

We conclude with further comments on the relationship with [6]. Their main result characterizes the class of distributions that admit a decomposition which is both learnable and sufficient for prediction. They dub these processes "asymptotically reverse mixing." In particular, they prove that, for every such process μ, the decomposition of μ induced by the future tail is learnable and sufficient to prediction. In our Example 3.3, the tail decomposition equals the ergodic decomposition, and, as we have shown, is not learnable. This shows that stationary processes needs not be asymptotic reverse mixing. On the other hand, the class of asymptotically reverse mixing processes contains non-stationary processes. For example, the Dirac atomic measure δ_ω is asymptotically reverse mixing for every realization $\omega \in \Delta(\Omega)$.

6. Extensions

In this section we discuss to what extent the theorems and tools of this paper extend to a larger class of process. In the process, this sheds further light on the assumptions made in our work.

[6] It is taken for granted in the first sentence of JKS's proof of Their Theorem 1

6.1. **Infinite set of outcomes.** The definitions of merging and weak merging can be extended to the case in which the outcome set A is a compact metric space[7]: Let ϕ be the Prohorov Metric over $\Delta(A)$. Say that the belief $\tilde{\mu} \in \Delta(A^{\mathbb{N}})$ *merges* to $\mu \in \Delta(A^{\mathbb{N}})$ if

$$\phi\left(\mu(\cdot|a_0,\dots,a_{n-1}),\tilde{\mu}(\cdot|a_0,\dots,a_{n-1})\right) \xrightarrow[n\to\infty]{} 0$$

for μ-almost every realization $\omega = (a_0, a_1, \dots) \in A^{\mathbb{N}}$ and that $\tilde{\mu}$ *weakly merges* to μ if the limit holds in strong Cesaro sense. Theorem 3.1 extends to the case of an infinite set A of outcomes. However, Theorem 5.2 does not hold in this case. We used the finiteness in the proof when we appealed to the equality of the past and future tails of the process. The following example shows the problem where A is infinite:

Example 6.1. Let $A = \{0,1\}^{\mathbb{N}}$ equipped with the standard Borel structure. Thus an element $a \in A$ is given by $a = (a[0], a[1], \dots)$ where $a[k] \in \{0,1\}$ for every $k \in \mathbb{N}$. Let μ be the belief over $A^{\mathbb{Z}}$ such that $\{\zeta_n[0]\}_{n\in\mathbb{Z}}$ are i.i.d. fair coin tosses and $\zeta_n[k] = \zeta_{n-k}[0]$ for every $k \geq 1$. Note that in this case $\overrightarrow{\mathcal{T}} = \mathcal{B}$ (so the future tail contains the entire history of the process) while $\overleftarrow{\mathcal{T}} = \mathcal{R}$ (the past tail is empty). The tail decomposition in this case will be the Dirac decomposition. However, this decomposition is not learnable: an agent who predict according to μ will at every period n will be completely in the dark about $\zeta_{n+1}[0]$. ▲

6.2. **Relaxing stationarity.** As we have argued earlier, stationary beliefs are useful to model situations where there is nothing remarkable about the point in time in which the agent started to keep track of the processes (so other agents who start observing the process at different times have the same beliefs) and that the agent is a passive observer who has no impact on the process itself. The first assumption is rather strong, and can be somewhat relaxed. In particular, consider a belief that is the posterior of some stationary prior conditioned on the occurrence of some event. (A similar situation is an agent who observes a finite state markov process that starts at a given state rather than the stationary distribution.) Let us say that a belief $\nu \in A^{\mathbb{N}}$ is *conditionally stationary* if there exists some stationary belief μ such that $\nu = \mu(\cdot|B)$ for some Borel subset B of $A^{\mathbb{N}}$ with $\mu(B) > 0$. While such processes are not stationary, they still admits an ergodic decomposition. they exhibit the same tail behavior of stationary processes. In particular, our theorems extend to such processes. The obvious details are omitted.

References

[1] P. Berti and P. Rigo. 0-1 laws for regular conditional distributions. *The Annals of Probability*, 35:649–662, 2007.

[7] Also for the case that A is a separable metric space, but then there are several possible non-equivalent definitions [4]

[2] David Blackwell and Lester Dubins. Merging of opinions with increasing information. *Ann. Math. Statist.*, 33:882–886, 1962.

[3] Thomas M Cover. Open problems in information theory. In *1975 IEEE Joint Workshop on Information Theory*, pages 35–36, 1975.

[4] Anthony D'Aristotile, Persi Diaconis, and David Freedman. On merging of probabilities. *Sankhyā: The Indian Journal of Statistics, Series A*, pages 363–380, 1988.

[5] AP Dawid. The Well-Calibrated Bayesian. *Journal of the American Statistical Association*, 77(379):605–610, 1982.

[6] Matthew O. Jackson, Ehud Kalai, and Rann Smorodinsky. Bayesian representation of stochastic processes under learning: de Finetti revisited. *Econometrica*, 67:875–893, 1999.

[7] E. Kalai and E. Lehrer. Rational learning leads to Nash equilibrium. *Econometrica*, 61:1019–1045, 1993.

[8] E. Kalai, E. Lehrer, and R. Smorodinsky. Calibrated Forecasting and Merging. *Games and Economic Behavior*, 29(1):151–159, 1999.

[9] Ehud Kalai and Ehud Lehrer. Weak and Strong Merging of Opinions. *Journal of Mathematical Economics*, 23:73–86, 1994.

[10] O. Kallenberg. *Foundations of Modern Probability*. Springer-Verlag, New York, second edition, 2002.

[11] E. Lehrer and R. Smorodinsky. Compatible measures and merging. *Mathematics of Operations Research*, pages 697–706, 1996.

[12] E. Lehrer and R. Smorodinsky. Relative entropy in sequential decision problems. *Journal of Mathematical Economics*, 33:425–439, 2000.

[13] Gusztáv Morvai and Benjamin Weiss. Forward estimation for ergodic time series. In *Annales de l'Institut Henri Poincare (B) Probability and Statistics*, volume 41, pages 859–870. Elsevier, 2005.

[14] Boris Yakovlevich Ryabko. Prediction of random sequences and universal coding. *Problemy Peredachi Informatsii*, 24(2):3–14, 1988.

[15] Benjamin Weiss. *Single Orbit Dynamics*. AMS Bookstore, 2000.

[16] H.V. Weizsäcker. Some reflections on and experiences with splifs. *Lecture Notes-Monograph Series*, pages 391–399, 1996.

Kellogg School of Management, Northwestern University

E-mail address: al-najjar@kellogg.northwestern.edu

School of Mathematics, Tel Aviv University and Kellogg School of Management, Northwestern University

E-mail address: erans@post.tau.ac.il